

Limits of kernel operators and the spectral regularity lemma

BALÁZS SZEGEDY

March 30, 2010

Abstract

We study the spectral aspects of the graph limit theory. We give a description of graphon convergence in terms of convergence of eigenvalues and eigenspaces. Along these lines we prove a spectral version of the strong regularity lemma. Using spectral methods we investigate group actions on graphons. As an application we show that the set of isometry invariant graphons on the sphere is closed in terms of graph convergence however the analogous statement does not hold for the circle. This fact is rooted in the representation theory of the orthogonal group.

Contents

1	Introduction	1
1.1	Hilbert-Schmidt kernel operators	3
1.2	The cut norm	4
1.3	Graph limits	6
1.4	Convergence in cut norm	7
1.5	Spectral Regularity lemma	9
1.6	Eigenspace convergence	10
1.7	Regularization and group actions	12
1.8	Weakly random group actions	14
1.9	Sphere vs. circle	15

1 Introduction

The so called graph limit theory (see [7],[10],[5],[2],[11],[12],[13],[14]) is a type of calculus developed on the completion of the set of finite graphs. As it was proved in [10], the elements of the completion can be represented by two variable symmetric functions $W : [0, 1]^2 \rightarrow [0, 1]$. It is very natural to interpret such functions as self adjoint integral kernel operators on $L_2([0, 1])$. Classical

theory says that every such operator has a spectral decomposition converging in L_2 . In this paper we focus on the spectral aspects of the graph limit theory. In the course of this investigation various interesting topics come up.

It was proved in [11] that the graph limit space is compact in the topology generated by a distance (called δ_\square) derived from the well known cut norm. This compactness implies a strong form [9] of Szemerédi's regularity lemma [15]. We give a new interpretation of cut norm convergence and δ_\square convergence in terms of spectral decompositions. Roughly speaking we prove that a sequence is convergent if and only if the eigenvalue sequences and the eigenspace structures converge in a rather strong way.

As a consequence we obtain a spectral form of the strong regularity lemma which can be regarded as a generalization of the strong regularity lemma by Alon, Fischer, Krivelevich and Szegedy [9]. We mention that numerous spectral aspects of the regularity lemma were studied by several authors. The closest approach to ours is by Frieze and Kannan [8]. An advantage of this type of regularization is that it is invariant under the symmetry group of the graph or graphon.

Using this fact we show that graphons or graphs can be regularized in a way that the structured part, which is a step function with bounded number of steps, is approximatively invariant under every automorphism of the graphon or graph. We call this statement the “symmetry preserving regularity lemma”. A symmetry preserving removal lemma was proved in [16]

If a unitary group action G on $L_2([0, 1])$ stabilizes a given graphon then the eigenspaces are also invariant under G . In particular they define finite dimensional representations of G . This creates an interesting connection between regularization and representation theory. In [6] Gowers proved that if in a finite group the dimension of the minimal non trivial irreducible representation is sufficiently big then its Cayley graphs are all arbitrarily quasi random. It is not hard to generalize this result for graphons with unitary group actions (see corollary 1.3). In the infinite case however a new interesting phenomenon appears. Let (V, μ) be a probability space. If a unitary action of G on $L_2(V, \mu)$ satisfies the condition that for every k there is a finite dimensional subspace of $L_2(V, \mu)$ containing all the G invariant subspaces of dimension at most k then we say the G acts weakly random. It turns out that graphons invariant under a weakly random action behave in a more controlled way. For example we prove that they form a closed set in the cut norm and so in the δ_\square distance. In particular the set of such graphons has a graph theoretic characterization using inequalities in subgraph densities.

Quite surprisingly, the circle behaves very different from the higher dimensional spheres. The set of isometry invariant graphons on the spheres of dimension ≥ 2 is closed in the δ_\square -norm and so it has a “graph theoretic characterization”. On the other hand isometry invariant graphons on the circle can have limits which can only be defined on the torus (or some other compact abelian group). This fundamental difference is coming from the fact that the action of O_{n+1} on $L_2(S_n)$ is weakly

random if and only if $n \geq 3$.

1.1 Hilbert-Schmidt kernel operators

Let us fix a separable probability space (V, μ) . Let \mathcal{H} denote the complex Hilbert space $L_2(V, \mu)$ with scalar product $(f, g) = \mathbb{E}_v(\overline{f(v)}g(v))$. The elements of \mathcal{H} are measurable functions so it makes sense to talk about their L_∞ (or L_1) norms even though the L_∞ norm might be infinite.

Recall that a sequence $\{f_i\}_{i=1}^\infty$ is called weakly convergent if $\{(f_i, g)\}_{i=1}^\infty$ is convergent for every $g \in \mathcal{H}$. It follows from the principle of uniform boundedness that weakly convergent sequences are bounded. Every weakly convergent sequence has a limit in \mathcal{H} . It is easy to see that every bounded sequence has a weakly convergent subsequence. It is known that convex bounded closed sets are weakly compact. For example the closed unit ball in the L_∞ norm is weakly compact.

Lemma 1.1 *Let $\{f_i\}$ be a weakly convergent sequence in \mathcal{H} with limit f such that $\lim_{i \rightarrow \infty} \|f_i\|_2 = \|f\|_2$ then $\{f_i\}_{i=1}^\infty$ converges to f in the L_2 norm.*

Proof. We have $\lim_{i \rightarrow \infty} \|f_i - f\|_2 = \lim_{i \rightarrow \infty} (f - f_i, f - f_i) = \lim_{i \rightarrow \infty} \|f\|_2 + \|f_i\|_2 - (f, f_i) - (f_i, f) = 0$.

A function $M : V \times V \rightarrow \mathbb{C}$ is called a Hilbert-Schmidt kernel operator if $M \in L_2(V \times V, \nu \times \nu)$. This is equivalent with saying that $M \in \mathcal{H} \otimes \mathcal{H}^*$. The operator M acts on \mathcal{H} by $f \mapsto Mf$ where $Mf(x) = \mathbb{E}_y(M(x, y)f(y))$. The image space $\text{ran}(M)$ of M is the Hilbert space generated by the functions $\{Mf | f \in \mathcal{H}\}$.

We will use the notion of weak convergence of kernel operators. A sequence of Hilbert-Schmidt kernel operators $\{H_i\}_{i=1}^\infty$ is called weakly convergent if they are weakly convergent in the Hilbert space $\mathcal{H} \otimes \mathcal{H}^* = L_2(V \times V)$. It is easy to see that if $\{\|H_i\|_2\}_{i=1}^\infty$ is a bounded sequence then $\{H_i\}_{i=1}^\infty$ is weakly convergent if and only if the sequences $\{g^* H_i f\}_{i=1}^\infty$ are convergent for every pair $f, g \in L_2(V)$.

An important consequence of the Cauchy-Schwartz inequality is that

$$\|Mg\|_2 \leq \|M\|_2 \|g\|_2 \quad \text{and} \quad |f^* Mg| \leq \|f\|_2 \|g\|_2 \|M\|_2 \quad (1)$$

for every $f, g \in \mathcal{H}$ and $M \in \mathcal{H} \otimes \mathcal{H}^*$. We will need the next lemma.

Lemma 1.2 *Let $\{f_i\}_{i=1}^\infty$ be a weakly convergent sequence in \mathcal{H} with limit f . If $M \in \mathcal{H} \otimes \mathcal{H}^*$ then $\lim_{i=1}^\infty \|Mf_i - Mf\|_2 = 0$.*

Proof. We have that $\|f_i\|_2 \leq c$ for every i with some positive constant c . Let $\{b_i\}_{i=1}^\infty$ be an orthonormal basis in \mathcal{H} . Then $M = \sum_{i,j} \alpha_{i,j} b_i^* b_j$ where $\sum_{i,j} |\alpha_{i,j}|^2 = \|M\|_2$. For every $\epsilon > 0$ there is a number t such that $M_t = \sum_{1 \leq i,j \leq t} \alpha_{i,j} b_i^* b_j$ satisfies $\|M - M_t\|_2 \leq \epsilon$. We have that by (1) that

$$\|Mf_i - Mf\|_2 = \|(M - M_t)f_i + (M - M_t)f + M_tf_i - M_tf\|_2 \leq 2\epsilon c + \|M_tf_i - M_tf\|_2.$$

If i is big enough then

$$\|M_t f_i - M_t f\|_2^2 = \sum_{i=1}^t \left| \sum_{j=1}^t \alpha_{i,j} (b_j, f_i - f) \right|^2$$

is smaller than $(\epsilon c)^2$ and for such indices $\|M f_i - M f\|_2 \leq 3\epsilon c$. Applying it for every $\epsilon > 0$ the proof is complete.

This implies immediately the next lemma.

Lemma 1.3 *Let $\{f_i\}_{i=1}^\infty$ and $\{g_i\}_{i=1}^\infty$ be two weakly convergent sequences in \mathcal{H} with limits f and g . Let M be a Hilbert-Schmidt kernel operator. Then $\lim f_i^* M g_i = f^* M g$.*

Let $M \in \mathcal{H} \otimes \mathcal{H}^*$ be a self adjoint Hilbert-Schmidt kernel operator. It is well known that M has a spectral decomposition

$$M = \sum_{i=1}^\infty f_i f_i^* \lambda_i$$

where $\{f_i\}_{i=1}^\infty$ is an orthogonal system of unit length elements in \mathcal{H} and the number's $\lambda_i \in \mathbb{R}$ satisfy $\sum_{i=1}^\infty |\lambda_i|^2 = \|M\|_2$. The numbers $\{\lambda_i\}_{i=1}^\infty$ are the eigenvalues of M listed with multiplicities.

Definition 1.1 *Let $M : V \times V \rightarrow \mathbb{C}$ be a self adjoint kernel operator with spectral decomposition $M = \sum_i f_i f_i^* \lambda_i$ and $\lambda \geq 0$. Then we denote by $[M]_\lambda$ the sum $\sum_{\{i \mid |\lambda_i| > \lambda\}} f_i f_i^* \lambda_i$.*

It is easy to see that $[M]_\lambda$ does not depend on the concrete choice of the functions f_i even if there are multiple eigenvalues. A basis independent definition of $[M]_\lambda$ is $[M]_\lambda = \sum_{|\tau| > \lambda} \tau P_\tau$ where P_τ is the orthogonal projection to the eigenspace $W_\tau = \{f \mid Mf = \tau f\}$.

Since M is a measurable function on $V \times V$ we can talk about the L_∞ and L_1 norms of M . Kernel operators with finite L_∞ norms will have a special importance for us.

Lemma 1.4 *Let M be a self adjoint kernel operator with $\|M\|_\infty \leq 1$. Assume that $f \in \mathcal{H}$ satisfies $\|f\|_2 = 1$ and $Mf = \lambda f$ for some non zero number λ . Then $\|f\|_\infty \leq |\lambda|^{-1}$.*

Proof. By Cauchy-Schwartz we have that $|\lambda f(x)| = |Mf(x)| = |\mathbb{E}_y(M(x, y)f(y))| \leq \|f\|_2 = 1$.

The spectral radius is an important invariant of kernel operators. It is defined as

$$\text{rad}(M) = \sup_{\|f\|_2=1} \|Mf\|_2.$$

The spectral radius of a self adjoint kernel operator is the maximum of the absolute values of its eigenvalues.

1.2 The cut norm

We will use the cut norm on $\mathcal{H} \otimes \mathcal{H}^*$ defined by

$$\|M\|_\square = \sup_{\|f\|_\infty, \|g\|_\infty \leq 1} |f^* M g|.$$

where f and g ranges over all possible measurable functions on V with L_∞ norm at most 1. Note that there are several definitions of the $\|\cdot\|_\square$ -norm that equivalent up to constant multiples. For example in [10] we used $\sup_{S,T} |\int_{S \times T} M|$ where S, T runs through all pairs of measurable sets in V .

Lemma 1.5 $\|M\|_\square \leq \text{spec}(M)$.

Proof. Let f, g be arbitrary functions with $\|f\|_\infty \leq 1$ and $\|g\|_\infty \leq 1$. We have that $\|f\|_2 \leq 1$ and $\|g\|_2 \leq 1$. Then $\|Mg\|_2 \leq \text{rad}(M)$ and so by Cauchy-Schwartz $f^*Mg \leq \text{rad}(M)$.

This implies the next lemma.

Lemma 1.6 *If $M : V \times V \rightarrow \mathbb{C}$ is a self adjoint Hilbert-Schmidt operator and $\alpha > 0$ then $\|M - [M]_\alpha\|_\square \leq \alpha$.*

Lemma 1.7 *If a bounded sequence of kernel operators $\{M_i\}_{i=1}^\infty$ converges to M in the cut norm then it converges to M in the weak topology in $\mathcal{H} \otimes \mathcal{H}^*$. In particular if $f, g \in \mathcal{H}$ then $\lim_{i \rightarrow \infty} f^*M_i g = f^*Mg$.*

Proof. Let S be the set of finite linear combinations of operators of the form fg^* where $f, g \in L_\infty(V)$. It is classical that S is dense in the Hilbert space $\mathcal{H} \otimes \mathcal{H}^*$. The cut norm convergence implies that $\lim_{i \rightarrow \infty} (Q, M_i) = (Q, M)$ for every $Q \in S$. Since $\{M_i\}_{i=1}^\infty$ is a bounded sequence it has to be weakly convergent with limit M .

The previous lemma with lemma 1.1 implies the next corollary.

Corollary 1.1 *If a sequence $\{M_i\}$ in $\mathcal{H} \otimes \mathcal{H}^*$ converges to M in the cut norm and $\lim_{i \rightarrow \infty} \|M_i\|_2 = \|M\|_2$ then $\{M_i\}_{i=1}^\infty$ converges to M in L_2 .*

Let $\psi : V \rightarrow V$ be a measure preserving map. This means that ψ is measurable and $\mu(\psi^{-1}(A)) = \mu(A)$ for every measurable subset $A \subseteq V$. If $W : V \times V \rightarrow \mathbb{C}$ is a kernel operator then we define W^ψ by $W^\psi(x, y) = W(\psi(x), \psi(y))$.

Let $\|\cdot\|_n$ be one of the norms $\|\cdot\|_1, \|\cdot\|_2, \|\cdot\|_\square$ on $\mathcal{H} \otimes \mathcal{H}^*$. We define the distance $\delta_n(W_1, W_2)$ on $\mathcal{H} \otimes \mathcal{H}^*$ by

$$\inf_{\psi_1, \psi_2 : V \rightarrow V} \|W_1^{\psi_1} - W_2^{\psi_2}\|_n$$

where ψ_1 and ψ_2 ranges over all possible measure preserving maps on V .

It is easy to see (and was pointed out in several papers) that one of the maps (say ψ_2) can be omitted and the other one can be assumed to be invertible. This means

$$\delta_n(W_1, W_2) = \inf_{\psi} \|W_1^\psi - W_2\|_n$$

where ψ ranges through all invertible measure preserving maps. This fact together with lemma 1.1 implies the next lemma.

Lemma 1.8 *Let $\{M_i\}_{i=1}^\infty$ be a δ_\square -convergent sequence with limit M . If $\|M\|_2 = \lim_{i \rightarrow \infty} \|M_i\|_2$ then $\{M_i\}_{i=1}^\infty$ converges to M also in δ_1 .*

Let \mathcal{M} denote the set of self adjoint operators in $\mathcal{H} \otimes \mathcal{H}^*$ with L_∞ norm at most 1. Let \mathcal{X} be the space obtained from \mathcal{M} by identifying operators that are δ_\square distance 0 from each other. This way $(\mathcal{X}, \delta_\square)$ becomes a metric space. The next theorem follows from the results in [11]:

Theorem 1 *The metric space $(\mathcal{X}, \delta_\square)$ is compact.*

1.3 Graph limits

Let \mathcal{W}_0 denote the set of symmetric measurable functions $W : [0, 1]^2 \rightarrow [0, 1]$. If $G = (V, E)$ is a finite simple graph on the vertex set $[k] = \{1, 2, \dots, k\}$ then the homomorphism density of G in W is defined by

$$t(G, W) = \int_{x_1, x_2, \dots, x_n} \prod_{(i,j) \in E, i < j} W(x_i, x_j) dx_1 dx_2 \dots dx_n$$

where x_1, x_2, \dots, x_k are in $[0, 1]$.

The elements of \mathcal{W}_0 are also called graphons. Two graphons are equivalent if their δ_\square distance is 0. Let \mathcal{X}_0 denote the set of equivalence classes of graphons. The set $(\mathcal{X}_0, \delta_\square)$ is a compact topological space. We call \mathcal{X}_0 the graph limit space. The theory of graph limits is basically the calculus on the graph limit space.

A sequence $\{W_i\}_{i=1}^\infty$ in \mathcal{W}_0 is δ_\square -convergent if and only if $\{t(G, W_i)\}_{i=1}^\infty$ is convergent for every simple graph G . Two graphons W_1, W_2 are equivalent if $t(G, W_1) = t(G, W_2)$ for every simple graph G . This implies that homomorphism densities are well defined on the elements of \mathcal{X} . Let \mathcal{G} be the set of finite simple graphs.

Definition 1.2 *A map $p : \mathcal{X}_0 \rightarrow \mathbb{R}$ is called a graph polynomial if there are finitely many graphs G_1, G_2, \dots, G_n in \mathcal{G} and real numbers $\lambda_1, \lambda_2, \dots, \lambda_n$ such that $p(W) = \sum_{i=1}^n t(G_i, W) \lambda_i$ for every $W \in \mathcal{X}_0$.*

If $G \in \mathcal{G}$ is the disjoint union of G_1 and G_2 then $t(G, W) = t(G_1, W)t(G_2, W)$ for every $W \in \mathcal{X}_0$. It follows that polynomials are closed under multiplications and so they are forming a commutative algebra (containing the constant functions) of δ_\square continuous functions on \mathcal{X}_0 . Let $K(\mathcal{X}_0)$ denote this algebra. The Stone-Weierstrass theorem implies that every continuous function on \mathcal{X}_0 can be approximated in L_∞ by some polynomial function in $K(\mathcal{X}_0)$. The next lemma shows that closed subsets of \mathcal{X}_0 can be characterized through inequalities in subgraph densities.

Lemma 1.9 *A set $C \subseteq \mathcal{X}_0$ is a closed subset of \mathcal{X}_0 then there is a countable set of graph polynomials $\{p_i\}_{i=1}^\infty$ such that $C = \cap_{i=1}^\infty \{W | p_i(W) \geq 0\}$.*

Proof. Let $d(W) = \inf\{\delta_\square(W, W_2) | W_2 \in C\}$ be the distance function from C . Since d is a continuous function we can approximate it with arbitrary L_∞ precision by graph polynomials. Let p'_i be an $1/i$ approximation of d in $K(\mathcal{X}_0)$ and let $p_i = 1/i - p'_i$. It is clear that $\{p_i\}_{i=1}^\infty$ is an appropriate system of polynomials.

1.4 Convergence in cut norm

In this part we examine the relationship between spectral decompositions and convergence in cut norm. We use the notation from the previous chapter. Let $\{M_i\}_{i=1}^\infty$ be a sequence of self adjoint kernel operators in $\mathcal{H} \otimes \mathcal{H}^*$ with $\|M_i\|_\infty \leq 1$ such that they converge in the cut norm. Let M be the cut norm limit of $\{M_i\}_{i=1}^\infty$. Obviously M is a self adjoint kernel operator and satisfies $\|M\|_\infty \leq 1$. We will keep this notation for the rest of this chapter and we will prove statements on the properties of the sequence $\{M_i\}_{i=1}^\infty$.

It is not hard to see that cut norm convergence implies the following type of convergence of the spectrums. If for a kernel operator W with spectrum $\{\lambda_i\}_{i=1}^\infty$ we define the random variable $X(W)$ that takes the value λ_i with probability $\lambda_i^4 (\sum_i \lambda_i^4)^{-1}$ then the k -th moment of $X(W)$ is equal to $t(C_{4+k}, W)/t(C_4, W)$ where C_n is the cycle of length n . It follows that from [10] that $\{X(M_i)\}_{i=1}^\infty$ converges to $X(M)$ in the weak topology of probability distributions. The spectrum with multiplicities is fully decodable from $X(W)$ and $t(C_4, W)$ and so at the level of spectrums the cut norm convergence is fully described.

In the rest of the chapter we study joint convergence of the spectrum and the eigenspaces.

Lemma 1.10 *Let $\{f_i\}_{i=1}^\infty$ be a weakly convergent sequence in \mathcal{H} with limit f such that $\|f_i\|_2 = 1$ for every i and $M_i f_i = f_i \lambda_i$ where $\lim_{i \rightarrow \infty} \lambda_i = \lambda \neq 0$. Then $\{f_i\}_{i=1}^\infty$ converges in L_2 to f and $Mf = \lambda f$.*

Proof. Note that lemma 1.4 implies that $\|f_i\|_\infty \leq |\lambda_i|^{-1}$. By lemma 1.3 we have that $\lim_{i \rightarrow \infty} f_i^* M f_i = f^* M f$. On the other hand

$$|f_i^*(M - M_i)f_i| \leq |\lambda_i|^{-2} \|M - M_i\|_\square$$

and from $\lambda \neq 0$ we get

$$0 = \lim_{i \rightarrow \infty} |f_i^*(M - M_i)f_i| = \lim_{i \rightarrow \infty} |f_i^* M f_i - \lambda_i| = |f^* M f - \lambda|. \quad (2)$$

Using $\|f_i\|_\infty \leq |\lambda_i|^{-1}$ we get

$$\lim_{i \rightarrow \infty} |f_i^*(M - M_i)f_i| \leq \lim_{i \rightarrow \infty} \|f\|_\infty |\lambda_i|^{-1} \|M - M_i\|_\square = 0. \quad (3)$$

It follows by (3), lemma 1.3 and by (2) that

$$\lim_{i \rightarrow \infty} \lambda_i(f, f_i) = \lim_{i \rightarrow \infty} f^* M_i f_i = \lim_{i \rightarrow \infty} f^* M f_i = f^* M f = \lambda$$

and so $\lim_{i \rightarrow \infty} (f, f_i) = 1$. Since $\|f_i\|_2 = 1$ and $\|f\|_2 \leq 1$ this is only possible if $\|f\|_2 = 1$ and so by lemma 1.1 $\{f_i\}_{i=1}^\infty$ converges to f in L_2 .

Now we need to show that $Mf = \lambda f$. Let g be any element in \mathcal{H} with $(g, f) = 0$. We have that

$$\begin{aligned} g^* M_i f &= g^* M_i (f - f_i) + g^* M_i f_i = g^* M_i (f - f_i) + \lambda_i (g, f_i) \leq \\ &\leq \|g\|_2 \|M_i\|_2 \|f - f_i\|_2 + \lambda_i (g, f_i). \end{aligned}$$

It implies that $\lim_{i \rightarrow \infty} g^* M_i f = 0$. On the other hand by lemma 1.7 $\lim_{i \rightarrow \infty} g^* M_i f = g^* M f$. It follows that Mf is orthogonal to every function g which is orthogonal to f . It follows that f is an eigenvalue of M and by $f^* M f = \lambda$ and $\|f\|_2 = 1$ the proof is complete.

Lemma 1.11 *Let $\lambda > 0$ be a number such that $\{-\lambda, \lambda\} \cap \text{spec}(M) = \emptyset$. Then*

1. $\lim_{i \rightarrow \infty} \text{rk}([M_i]_\lambda) = \text{rk}([M]_\lambda)$,
2. $\lim_{i \rightarrow \infty} \|[M_i]_\lambda - [M]_\lambda\|_2 = 0$.

Proof. Assume that $M_i = \sum_{j=1}^\infty f_{i,j} f_{i,j}^* \lambda_{i,j}$ such that $\{|\lambda_{i,j}|\}_{j=1}^\infty$ is a decreasing sequence and the vectors $\{f_{i,j}\}_{\lambda_{i,j} \neq 0}$ are forming an orthonormal system. To keep the sequences $\{\lambda_{i,j}\}_{j=1}^\infty$ infinite we put an infinite number of 0's at the end if M_i has finite rank. (If $\lambda_{i,j} = 0$ then $f_{i,j}$ is an arbitrarily chosen function of unit length.)

First of all we prove that there is a subsequence $\{M_i\}_{i \in S}$ satisfying the condition of the lemma. By a standard argument we can choose a subsequence S such that $\{f_{i,j}\}_{i \in S}$ is weakly convergent for every fixed j and $\{\lambda_{i,j}\}_{i \in S}$ is convergent for every j . Let f_j be the weak limit of $\{f_{i,j}\}_{i \in S}$ and λ_j be the limit of $\{\lambda_{i,j}\}_{i \in S}$. Obviously we have that $\sum_{j=1}^\infty |\lambda_j|^2 \leq 1$ and $\{|\lambda_j|\}_{j=1}^\infty$ is a decreasing sequence. It follows that $|\lambda_j| \leq 1/\sqrt{j}$. First of all note that if $\lambda_j \neq 0$ then by lemma 1.10 $\lim_{i \rightarrow \infty} \|f_{i,j} - f_i\|_2 = 0$. It follows that if λ_{j_1} and λ_{j_2} are both non-zero then $(f_{j_1}, f_{j_2}) = 0$. In other words $\{f_j\}_{\{\lambda_j \neq 0\}}$ is an orthonormal system of functions. Let

$$M' = \sum_{\{j|\lambda_j \neq 0\}} f_j f_j^* \lambda_j.$$

First we claim that $M = M'$. For every natural number t we have that $\|M_i - M'\|_\square$ is at most

$$\left\| \sum_{j=1}^t (f_{i,j} f_{i,j}^* \lambda_{i,j} - f_j f_j^* \lambda_j) \right\|_\square + \left\| \sum_{j=t+1}^\infty f_{i,j} f_{i,j}^* \lambda_{i,j} \right\|_\square + \left\| \sum_{j=t+1}^\infty f_j f_j^* \lambda_j \right\|_\square.$$

The spectral radius of the sums in the last two terms is at most $1/\sqrt{t+1}$. It follows that if i is big enough then $\|M_i - M'\|_\square \leq 3/\sqrt{t+1}$. By letting t go to infinity we get that M_i converges to M' in the cut norm and so $M' = M$.

Let t be an integer greater than λ^{-2} . We have that $|\lambda_{i,j}| \leq \lambda$ whenever $j > t$. Since $\lambda, -\lambda$ are not eigenvalues of M we have that there is an index i_0 such that for $|\lambda_{i,j} - \lambda_j| \leq |\lambda - |\lambda_j||/2$

whenever $1 \leq j \leq t$ and $i > i_0$, $i \in S$. This means that for such indices $|\{j | |\lambda_{i,j}| > |\lambda|\}| = |\{j | \lambda_j > |\lambda|\}| = \text{rk}([M]_\lambda)$ showing that $\text{rk}([M_i]_\lambda) = \text{rk}([M]_\lambda)$.

Now we finish the general case by contradiction. If the first statement is not true then we can choose an infinite subsequence where $\text{rk}([M_i]_\lambda) \neq [M]_\lambda$. This is a contradiction since from such a subsequence we can not choose a sub sequence satisfying the first condition. If the second condition fails then we can choose an infinite subsequence for some ϵ such that $\|[M_i]_\lambda - [M]_\lambda\|_2 > \epsilon$. This is again a contradiction

1.5 Spectral Regularity lemma

Theorem 2 (Spectral regularity lemma) *For an arbitrarily decreasing function $F : \mathbb{R}^+ \times \mathbb{R}^+ \rightarrow \mathbb{R}^+$ and every $\epsilon > 0$ there is a constant $\delta > 0$ such that for every self adjoint kernel operator $M : V \times V \rightarrow \mathbb{C}$ with $\|M\|_\infty \leq 1$ on a separable probability space (V, μ) there is a real number $\lambda \geq \delta$ such that M has a decomposition $M = S + E + R$ with the following properties*

1. $S = [M]_\lambda$
2. $\|E\|_2 \leq \epsilon$
3. $\|R\|_\square \leq F(\lambda, \epsilon)$
4. $\|S + E\|_\infty \leq 1$
5. E and R are self adjoint.

Proof. We go by contradiction. Let $\epsilon > 0$ be a real number such that the theorem fails for ϵ . This means that there is a sequence of kernel operators $\{M_i\}_{i=1}^\infty$ with L_∞ norm at most 1 such that M_i does not have the desired decomposition for $\delta = 1/i$. We can assume without loss of generality that all the operators M_i are defined on the same standard probability space V . Also without loss of generality (by choosing a subsequence guaranteed by theorem 1) we can assume that $\{M_i\}_{i=1}^\infty$ is convergent in δ_\square and so there is a sequence of invertible measure preserving maps $\{\psi_i\}_{i=1}^\infty$ on V such that $\{M_i^{\psi_i}\}_{i=1}^\infty$ converges to M with $\|M\|_\infty \leq 1$ in the cut norm. Let $\lambda > 0$ be a number such that $\{\lambda, -\lambda\} \cap \text{spec} M = \emptyset$ and $\|[M]_\lambda - M\|_2 \leq \epsilon/3$. By lemma 1.11 there is an index i_0 such that for $i > i_0$ we have $\|[M_i^{\psi_i}]_\lambda - [M]_\lambda\|_2 \leq \epsilon/3$. This means that if $i > i_0$ the $\|M - [M_i^{\psi_i}]_\lambda\|_2 \leq 2\epsilon/3$. Let $E_i = M - [M_i^{\psi_i}]_\lambda$ and $R_i = M_i^{\psi_i} - M$. Now $M_i^{\psi_i} = [M_i^{\psi_i}]_\lambda + E_i + R_i$. Now since R_i converges to 0 in the cut norm it follows that there is an index $i_1 > \max(i_0, 1/\lambda)$ such that if $i > i_1$ then $\|R_i\|_\square < F(\lambda, \epsilon)$ satisfies the theorem with $1/i$. Applying ψ^{-1} to the decomposition of $M_i^{\psi_i}$ we get a contradiction.

Now let us assume that V is a finite probability space with uniform distribution. We can represent undirected graphs on the vertex set V by their adjacency matrices $G : V \times V \rightarrow \{0, 1\}$. More generally assume that $G : V \times V \rightarrow \mathbb{R}$ is a symmetric matrix. The automorphism group $\text{Aut}(G)$ is

the group of permutation matrices g satisfying $gGg^{-1} = G$. A matrix $G : V \times V \rightarrow \mathbb{R}$ is called a step function with n -steps if there is a partition $\mathcal{P} = \{P_1, P_2, \dots, P_n\}$ of V and an n by n matrix $T : [n]^2 \rightarrow \mathbb{R}$ such that $G(a, b) = T(i, j)$ whenever $a \in P_i$ and $b \in P_j$. We say that G is a balanced step function with n steps if $||P_i| - |P_j|| \leq 1$ for every i, j .

Remark 1.1 *Let $C \subseteq \mathcal{H}$ be a convex, L_∞ -bounded closed set which is invariant under measure preserving maps on V . Then similar regularity lemma holds for every M in C such that $S + E \in C$. In particular if C is the set of kernel operators taking values in $[0, 1]$ then M is a graphon and so is $S + E$. The proof is essentially the same.*

Lemma 1.12 (Eigenvector clustering) *If $\epsilon > 0$ and $G = \sum_{i=1}^k f_i f_i^* \lambda_i$ such that $\|f_i\|_\infty \leq m$ and $|\lambda_i| \leq m$ for every $1 \leq i \leq k$ then there is a step function T with at most $(20km^3/\epsilon)^k$ steps such that $\|T - G\|_\infty \leq \epsilon$.*

Proof. Let us consider the partition $V = \cup_{i=1}^t P_i$ according to the level sets of the function

$$v \mapsto (\lfloor f_i(v) \epsilon_1^{-1} \rfloor \epsilon_1)_{i=1}^k.$$

Here $t \leq (2m/\epsilon_1)^k$. If two element $v_1, v_2 \in P_i$ and $w_1, w_2 \in P_j$ then with a rough estimate $|T(v_1, w_1) - T(v_2, w_2)| \leq 10km^2\epsilon_1$. It follows that there is step function T with partition $\{P_i\}_{i=1}^t$ such that $\|T - G\|_\infty \leq 10km^2\epsilon_1$. If $\epsilon_1 \leq \epsilon/(10km^2)$ then T satisfies the condition of the lemma.

Using the fact that in theorem 2 the matrix $[M]_\lambda$ is invariant under the automorphisms of G and the previous lemma we obtain the next version of the classical graph regularity lemma.

Theorem 3 (Symmetry preserving regularity lemma) *For an arbitrarily decreasing function $F(\mathbb{R}^+, \mathbb{N}) \rightarrow \mathbb{R}^+$ and $\epsilon > 0$ there is constant n such that for every symmetric matrix $G : V \times V \rightarrow [-1, 1]$ there is decomposition $G = S + E + R$ such that*

1. S is a step function with $s \leq n$ steps
2. $\|gSg^{-1} - S\|_\infty \leq \epsilon$ for every $g \in \text{Aut}(G)$
3. $\|E\|_2 \leq \epsilon$
4. $\|R\|_\square \leq F(\epsilon, s)$.

At the cost of worsening the bound of n in terms of F and ϵ we can also assume that T is a balanced step function. However in this case the L_∞ error in $\|gSg^{-1} - S\|_\infty \leq \epsilon$ becomes an L_2 error.

1.6 Eigenspace convergence

Let \mathcal{M} be the set of kernel operators $M : V \times V \rightarrow \mathbb{C}$ with $\|M\|_\infty \leq 1$.

Proposition 1.1 *Let $\{M_i\}_{i=1}^\infty$ be a sequence in \mathcal{M} . Then the following two statements are equivalent.*

1. $\{M_i\}_{i=1}^\infty$ is convergent in the cut norm
2. there is a decreasing positive real sequence $\{\alpha_i\}_{i=1}^\infty$ with $\lim_{i \rightarrow \infty} \alpha_i = 0$ such that $\{[M_i]_{\alpha_j}\}_{i=1}^\infty$ is L_2 convergent for every j .

Furthermore in the second statement the cut norm limit M of $\{M_i\}_{i=1}^\infty$ can be computed as

$$M = \lim_{j \rightarrow \infty} \left(\lim_{i \rightarrow \infty} [M_i]_{\alpha_j} \right)$$

converging in L_2 .

Proof. To show that the first statement implies the second one let S be the set of eigenvalues of the cut norm limit M of $\{M_i\}_{i=1}^\infty$. Then by lemma 1.11 any sequence $\{\alpha_i\}_\infty$ avoiding the absolute values of the eigenvalues of M satisfies the convergence requirement second statement. The eigenvalues of M are forming a countable set and so we can choose $\{\alpha_i\}_{i=1}^\infty$ satisfying the required conditions.

Let $L_j = \lim_{i \rightarrow \infty} [M_i]_{\alpha_j}$ and $D_j = L_{j+1} - L_j$. The functions $D_{i,j} = [M_i]_{\alpha_{j+1}} - [M_i]_{\alpha_j}$ satisfy the following properties for every i

1. $D_{i,j} D_{i,k} = 0$ for every $i \neq k$,
2. $(D_{i,j}, D_{i,k}) = 0$ in $\mathcal{H} \otimes \mathcal{H}^*$,
3. the spectral radius of $D_{i,j}$ is at most α_j ,
4. $\sum_{j=1}^\infty \|D_{i,j}\|_2^2 = \|M_i\|_2^2 \leq 1$.

This means that the sequence $\{D_j\}_{j=1}^\infty$ satisfies the same properties. Using that last and the second property we have that $\sum_{j=1}^\infty D_j$ is convergent in the L_2 norm. Let us denote the limit by M . Our goal is to show that $\{M_i\}_{i=1}^\infty$ converges to M in the cut norm.

The first and third property implies that the spectral radius of $N_j = \sum_{k=j}^\infty D_k$ is at most α_j and so

$$\begin{aligned} \|M - M_i\|_\square &\leq \|(M - N_{j+1}) - [M_i]_{\alpha_j}\|_\square + \|M_i - [M_i]_{\alpha_j}\|_\square + \|N_{j+1}\|_\square \leq \\ &\leq \|L_j - [M_i]_{\alpha_j}\|_2 + 2\alpha_j. \end{aligned}$$

Now using the fact that $[M_i]_{\alpha_j}$ converges to L_j we get that if i is big enough then $\|M - M_i\|_\square \leq 3\alpha_j$. Since $\{\alpha_j\}_{j=1}^\infty$ converges to 0 the proof is complete.

Lemma 1.13 *If $\{M_i\}_{i=1}^\infty$ in \mathcal{M} converges to M in the δ_\square distance then for every $\lambda > 0$ we have $\limsup_{i \rightarrow \infty} \|[M]_\lambda\|_2 \leq \|M\|_2$.*

Proof. We can choose a sequence of invertible measure preserving maps $\{\psi_i\}_{i=1}^\infty$ such that $\{M_i^{\psi_i}\}_{i=1}^\infty$ is convergent in the cut norm. By lemma 1.11 there is a value $0 < \alpha < \lambda$ such that $\lim_{i \rightarrow \infty} [M_i]_\alpha = [M]_\alpha$ in L_2 . This means that $\limsup_{i \rightarrow \infty} \|[M_i]_\lambda\|_2 \leq \lim_{i \rightarrow \infty} \|[M_i]_\alpha\|_2 \leq \|M_2\|_2$.

Proposition 1.2 *Let $\{M_i\}_{i=1}^\infty$ be a sequence in \mathcal{M} . Then the following two statements are equivalent.*

1. $\{M_i\}_{i=1}^\infty$ is convergent in the δ_\square distance
2. there is a decreasing positive real sequence $\{\alpha_i\}_{i=1}^\infty$ with $\lim_{i \rightarrow \infty} \alpha_i = 0$ such that $\{[M_i]_{\alpha_j}\}_{i=1}^\infty$ is δ_1 convergent for every j .

Furthermore in the second statement the cut norm limit M of $\{M_i\}_{i=1}^\infty$ can be computed as

$$M = \lim_{j \rightarrow \infty} \left(\lim_{i \rightarrow \infty} [M_i]_{\alpha_j} \right)$$

converging in δ_1 .

Proof. To see that the first statement implies the second choose a sequence of invertible measure preserving transformations $\{\psi_i\}_{i=1}^\infty$ such that $\{M_i^{\psi_i}\}_{i=1}^\infty$ converges in the cut norm. Then proposition 1.1 shows the second statement.

We show that the second statement implies the first. $L_j = \lim_{i \rightarrow \infty} [M_i]_{\alpha_j}$. Observe that by lemma 1.6 $\|M_i - [M_i]_{\alpha_j}\|_\square \leq \alpha_j$ and so $\delta_\square(M_i, [M_i]_{\alpha_j}) \leq \alpha_j$. If i is big enough that $\delta_\square(M_i, L_j) \leq 2\alpha_j$. This means that $\{M_i\}_{i=1}^\infty$ is a δ_\square Cauchy sequence which shows the first statement. It also shows that $\{L_j\}_{j=1}^\infty$ converges to the cut norm limit M of $\{M_i\}_{i=1}^\infty$ in the δ_\square -distance. It remains to show that this convergence is also true in the δ_1 metric. This follows from the fact that by lemma 1.13 $\|L_j\|_2 \leq \|M\|_2$ and so lemma 1.8 completes the proof.

1.7 Regularization and group actions

An advantage of the spectral regularity lemma is that it is invariant under the symmetries of kernel operators. To be more precise let G be a group of unitary operators on $\mathcal{H} = L_2(V)$. Then there is a natural induced action of G on $\mathcal{H} \otimes \mathcal{H}^*$ and in particular on the set of self adjoint Hilbert-Schmidt kernel operators. This action satisfies $H^\alpha f^\alpha = (Hf)^\alpha$ and $(f^\alpha)^* H^\alpha = (f^* H)^\alpha$. If $H = \sum_i f_i f_i^* \lambda_i$ is a spectral decomposition of H then $H^\alpha = \sum_i f_i^\alpha (f_i^\alpha)^* \lambda_i$. The next well known lemma is trivial from the previous remarks.

Lemma 1.14 *If a group G of unitary operators on $L_2(V)$ stabilizes a self adjoint Hilbert-Schmidt kernel operator H ($H^\alpha = H$ for every $\alpha \in G$) then G stabilizes all the operators $[H]_\lambda$ for $\lambda \geq 0$. Furthermore the eigenspaces W_{λ_i} of H are G invariant spaces. In particular $\text{ran}(H)$ is G invariant.*

A typical example for a kernel operator stabilized by a group action is a Cayley graphon. Let G be a compact Hausdorff topological group with normalized Haar measure μ . Let f be a Borel measurable function $f : G \rightarrow \mathbb{C}$. Let $M : G \times G \rightarrow \mathbb{C}$ be the kernel operator defined by $M(x, y) = f(x^{-1}y)$. It is easy to see that the left action of G on itself induces a unitary group action of G on $L_2(G, \mu)$ and it stabilizes M . If f has the property that $f(g^{-1}) = \overline{f(g)}$ then the corresponding Cayley graphon is self adjoint. The next corollary of lemma 1.14 creates a connection between quasi randomness and representation theory.

Corollary 1.2 (Quasirandom action) *If a group G of unitary operators on $L_2(V)$ stabilizes a self adjoint Hilbert-Schmidt operator H with $\|H\|_2 \leq 1$ then the spectral radius (and so the cut norm) of H is at most $1/\sqrt{d}$ where d is the smallest dimension of a G invariant subspace in $\text{ran}(H)$.*

Proof. Lemma 1.14 implies that every eigenvalue of H has multiplicity at least d . Since the sum of the squares of the eigenvalues is at most 1 we get that $\lambda^2 d < 1$ is satisfied by every eigenvalue λ .

Corollary 1.3 (Quasirandom action II.) *Let G be a compact Hausdorff topological group with normalized Haar measure μ . Let $K \subseteq G$ be a closed subgroup and V be the left coset space $\{gK | g \in G\}$. Let $\mathcal{H} = L_2(V, \mu)$ and d be the degree of the smallest non trivial representation of G which appears in the induced action of G on \mathcal{H} . Then every G invariant self adjoint Hilbert-Schmidt kernel operator H with $\|H\|_2 \leq 1$ satisfies $\|H - p\|_{\square} \leq 1/\sqrt{d}$ where p is the constant function on $V \times V$ with value $\alpha = \int_{x,y} H(x, y) dx dy$.*

Proof. Let \mathcal{H}^0 denote the orthogonal space of the constant 1 function on V . The smallest finite dimensional G invariant subspace in \mathcal{H}^0 has dimension at least d . It is easy to see that $\text{ran}(H - p) \subseteq \mathcal{H}^0$. Then corollary 1.2 finishes the proof.

We demonstrate the usefulness of this simple fact on the next example. Let S_n denote the n -dimensional sphere with the isometry invariant probability measure. We call a graphon $W : S_n \times S_n \rightarrow [0, 1]$ isometry invariant if W is invariant under the induced action of the orthogonal group O_{n+1} on S_n . Note that a graphon W is isometry invariant if and only if the value $W(x, y)$ depends only on the distance of x and y .

The next proposition says that on a very high dimensional sphere every isometry invariant graphon is very close to being quasirandom.

Proposition 1.3 *If W is an isometry invariant graphon with edge density p on S_n then $\|W - p\|_{\square} \leq 1/\sqrt{n+1}$.*

Proof. The smallest non trivial representation of the orthogonal group O_{n+1} which appears on $L_2(S_n)$ has dimension $n+1$. Then corollary 1.3 completes the proof.

1.8 Weakly random group actions

Definition 1.3 Let $G \subseteq U(L_2(V))$ be a unitary operator group. We denote by $\mathcal{I}(G, d)$ the set of G invariant self adjoint kernel operators with L_∞ norm at most d .

Lemma 1.15 The set $\mathcal{I}(G, d)$ is closed under weak convergence.

Proof. Let $\{H_i\}_{i=1}^\infty$ be a weakly convergent sequence of kernel operators in $\mathcal{I}(G, d)$ and let H be the weak limit. For every two functions $f, g \in L_2(V)$ we have that $g * H^\alpha f = \lim g * H_i^\alpha f = \lim g * H_i f = g * H f$. This means that $H^\alpha = H$.

We will need the following lemma about weak convergence.

Lemma 1.16 Let $C \subseteq L_2(V)$ be a compact set and $\{H_i\}_{i=1}^\infty$ be a sequence of kernel operators, with uniformly bounded L_2 norms, weakly converging to the 0 function. Then

$$\lim_{i \rightarrow \infty} \max_{f, g \in C} \|g * H_i f\|_2 = 0.$$

Proof. Using compactness of C we can choose sequences $\{g_i\}_{i=1}^\infty$ and $\{f_i\}_{i=1}^\infty$ in C such that $\|g_i * H_i f_i\|_2 = \max_{f, g \in C} \|g * H_i f\|_2 = m_i$. Assume that $m = \limsup_{i \rightarrow \infty} m_i > 0$. Then by choosing a subsequence we can assume that $m = \lim_{i \rightarrow \infty} m_i$. Furthermore by compactness of C we can assume by choosing a subsequence that $\lim_{i \rightarrow \infty} f_i = f$ and $\lim_{i \rightarrow \infty} g_i = g$ where the convergence is in the L_2 norm. Now using the fact that the L_2 norms of H_i are bounded we obtain that $\lim_{i \rightarrow \infty} \|g * H_i f\|_2 = m > 0$ which is a contradiction.

Definition 1.4 Let G be a group of unitary operators on $L_2(V)$. We say that G acts weakly random if for every natural number n the space U_n generated by all G invariant subspaces in $L_2(V)$ of dimension at most n is finite dimensional.

Next theorem shows a surprising graph theoretic aspect of weakly random group actions.

Theorem 4 Let G be a weakly random operator group and let $\mathcal{I}(G, d)$ denote the set of self adjoint integral kernel operators $M : V \times V \rightarrow \mathbb{C}$ that are invariant under G and have L_∞ norm at most $d < \infty$. Then weak convergence on $\mathcal{I}(G, d)$ coincides with convergence in the cut norm. In particular $\mathcal{I}(G, d)$ is cut norm compact.

Proof. Let $\{H_i\}_{i=1}^\infty$ be a weakly convergent sequence in $\mathcal{I}(G, d)$. By subtracting (from every term) the weak limit H (which is also in $\mathcal{I}(G, d)$ by lemma 1.15) we get a sequence in $\mathcal{I}(G, 2d)$ which converges to 0 weakly. This means that without loss of generality we can assume that the weak limit of $\{H_i\}_{i=1}^\infty$ is the 0 function. Let us choose an arbitrary real number $\epsilon > 0$. Let U be the space generated by all G invariant subspaces of dimension at most d^2/ϵ^2 and let U^0 be the unit ball in U . Using lemma 1.14 we get that every normalized eigenvector of H_i corresponding to an eigenvalue

of absolute value bigger than ϵ is in U^0 . On the other hand, using the compactness of U_0 and lemma 1.16 we get that there is an index j such that if $i > j$ then $\max_{f \in U^0} \|f^* H_i f\|_2 < \epsilon$. This means that if $i > j$ then the spectral radius of H_i is at most ϵ and so $\|H_i\|_{\square} \leq \epsilon$.

1.9 Sphere vs. circle

Let $f : [-1, 1] \rightarrow \mathbb{R}$ be a bounded measurable function. We denote by $S(n, f)$ the graphon defined on the unit sphere $S_n = \{x \mid x \in \mathbb{R}^{n+1}, \|x\|_2 = 1\}$ with the uniform distribution such that $w(x, y) = f(xy)$ where xy is the usual scalar product.

The underlying topological space (in the sense of [13]) of $S(n, f)$ is either the sphere S_n or just one point if f is constant. In this part we point out that the case $n = 1$ is very different for $n > 1$. Let \mathcal{S}_n^0 denote the subset in $S(n, f)$ where $0 \leq f(x) \leq 1$ for every $-1 \leq x \leq 1$.

Proposition 1.4 *Let $n \geq 2$. Then the set \mathcal{S}_n^0 is compact in the cut norm.*

Proof. For the first part let O_{n+1} be the orthogonal group acting on S_n . The induced action of O_n on $L_2(S_n)$ is defined by $f^\alpha(x) = f(x^\alpha)$ where $\alpha \in O_{n+1}$. It is clear that spherical graphons of dimension n are invariant under this action. The representation theory of O_{n+1} on $L_2(S_n)$ is a classical theory. It acts weakly randomly which proves the first part.

Corollary 1.4 *Let $n \geq 2$. Then the set \mathcal{S}_n^0 is closed in δ_{\square} and so it is characterizable by inequalities in subgraph densities.*

Proof. If $\{M_i\}_{i=1}^\infty$ in \mathcal{S}_n^0 is δ_{\square} convergent then we can choose a weakly convergent subsequence. By proposition 1.4 this subsequence is cut norm convergent. The limit M is in \mathcal{S}_n^0 and it has to coincide with the δ_{\square} limit of $\{M_i\}_{i=1}^\infty$.

Proposition 1.5 *If $n = 1$ then the set \mathcal{S}_n^0 is not compact in the cut norm.*

Proof. Let us define the graphon W on the circle S_1 by $W(x, y) = 1$ if $xy > 0$ and $W(x, y) = 0$ if $xy < 0$. Let us represent S_1 and the abelian group $A = \mathbb{R}/\mathbb{Z}$. It is easy to see a well known that for every fixed k the map $\psi_k : a \mapsto ka$ is a measure preserving map on A . This means that $\delta_{\square}(W^{\psi_k}, W) = 0$ for every k . We show that the sequence $\{W^{\psi_k}\}_{k=1}^\infty$ does not have a cut norm convergent sub sequence. Assume by contradiction the $\{W^{\psi_{k_i}}\}_{i=1}^\infty$ is convergent in the cut norm. Then the limit L has δ_{\square} distance 0 from W and thus $\|L\|_2 = \|W\|_2$. By lemma 1.8 this means that $\{W^{\psi_{k_i}}\}_{i=1}^\infty$ is convergent in L_2 . It is easy to see that this is not the case.

Without proof we mention that the set \mathcal{S}_1^0 is not even closed in δ_{\square} . There are examples where a sequence of graphons in \mathcal{S}_1^0 converges to a graphon whose underlying topological space is the torus. Such a graphon can't be represented on the circle.

Motivated by the above results it is natural to introduce the following notion.

Definition 1.5 Let (V, \mathcal{A}, μ) be a probability space with σ -algebra \mathcal{A} and measure μ . Let $\mathcal{B} \subseteq \mathcal{A} \times \mathcal{A}$ be a sub σ -algebra on the product space $V \times V$. We call \mathcal{B} weakly random if in the set of functions $\{M | M \in L_\infty(\mathcal{B}), \|M\|_\infty \leq 1\}$ weak convergence implies cut norm convergence.

Proposition 1.4 says that if V is the sphere S_2 and \mathcal{B} consists of those Borel measurable sets that are invariant under the diagonal action of O_3 on $V \times V$ then \mathcal{B} is weakly random.

Question 1 Is there any characterization of weakly random σ -algebras?

Another interesting topic is to understand when $\{M | M \in L_\infty(\mathcal{B}), \|M\|_\infty \leq 1\}$ is closed in the δ_\square metric. Weakly randomness implies this but the other direction is not true.

Balázs Szegedy University of Toronto, Department of Mathematics, St George St. 40, Toronto, ON, M5R 2E4, Canada

References

- [1] G. Elek, B. Szegedy: *A measure-theoretic approach to the theory of dense hypergraphs*, preprint
- [2] C. Borgs, J. Chayes, L. Lovász: Moments of Two-Variable Functions and the Uniqueness of Graph Limits, *Geometric and Functional Analysis* (to appear)
- [3] C. Borgs, J. Chayes, L. Lovász, V.T. Sós, K. Vesztergombi: Counting graph homomorphisms, in: *Topics in Discrete Mathematics* (ed. M. Klazar, J. Kratochvil, M. Loeb, J. Matoušek, R. Thomas, P. Valtr), Springer (2006), 315–371.
- [4] C. Borgs, J.T. Chayes, L. Lovász, V.T. Sós, B. Szegedy and K. Vesztergombi: Graph Limits and Parameter Testing, Proc. 38th Annual ACM Symp. on Theory of Computing 2006, 261–270.
- [5] C. Borgs, J.T. Chayes, L. Lovász, V.T. Sós, and K. Vesztergombi: Convergent Graph Sequences I: Subgraph frequencies, metric properties, and testing, *Advances in Math.* (2008), 10.1016/j.aim.2008.07.008.
- [6] W.T. Gowers, *Quasirandom groups* preprint
- [7] M. Freedman, L. Lovász, A. Schrijver: Reflection positivity, rank connectivity, and homomorphisms of graphs, *J. Amer. Math. Soc.* **20** (2007), 37–51.
- [8] A. Frieze and R. Kannan: Quick approximation to matrices and applications, *Combinatorica* **19**, 175–220.
- [9] N. Alon, E. Fischer, M. Krivelevich, M. Szegedy: Efficient testing of large graphs, *Combinatorica* **20** (2000), 451–476

- [10] L. Lovász, B. Szegedy: Limits of dense graph sequences, *J. Comb. Theory B* **96** (2006), 933–957.
- [11] L. Lovász and B. Szegedy: Szemerédi’s Lemma for the analyst, *Geom. Func. Anal.* **17** (2007), 252–270.
- [12] L. Lovász and B. Szegedy: Finitely forcible graphons (submitted)
- [13] L. Lovász and B. Szegedy: Regularity partitions and the topology of graphons (submitted)
- [14] L. Lovász and B. Szegedy: Testing properties of graphs and functions (submitted)
- [15] E. Szemerédi: Regular partitions of graphs, *Colloque Inter. CNRS* (J.-C. Bermond, J.-C. Fournier, M. Las Vergnas and D. Sotteau, eds.) (1978) 399–401.
- [16] B. Szegedy: The Symmetry Preserving Removal Lemma (submitted)